ARTICLE

Tripled yield in direct–drive laser fusion through statistical modelling

V. Gopalaswamy[1,2*], R. Betti[1,2,3], J. P. Knauer[1], N. Luciani[1,2,4], D. Patel[1,2], K. M. Woo[1,3], A. Bose[1,5], I. V. Igumenshchev[1], E. M. Campbell[1], K. S. Anderson[1], K. A. Bauer[1], M. J. Bonino[1], D. Cao[1], A. R. Christopherson[1,2], G. W. Collins[1], T. J. B. Collins[1], J. R. Davies[1], J. A. Delettrez[1], D. H. Edgell[1], R. Epstein[1], C. J. Forrest[1], D. H. Froula[1], V. Y. Glebov[1], V. N. Goncharov[1], D. R. Harding[1], S. X. Hu[1], D. W. Jacobs-Perkins[1], R. T. Janezic[1], J. H. Kelly[1], O. M. Mannion[1,3], A. Maximov[1,2], F. J. Marshall[1], D. T. Michel[1], S. Miller[1,2], S. F. B. Morse[1], J. Palastro[1], J. Peebles[1], P. B. Radha[1], S. P. Regan[1], S. Sampat[1], T. C. Sangster[1], A. B. Sefkow[1], W. Seka[1], R. C. Shah[1], W. T. Shmyada[1], A. Shvydky[1], C. Stoeckl[1], A. A. Solodov[1], W. Theobald[1], J. D. Zuegel[1], M. Gatu Johnson[5], R. D. Petrasso[5], C. K. Li[5] & J. A. Frenje[5]

Focusing laser light onto a very small target can produce the conditions for laboratory-scale nuclear fusion of hydrogen isotopes. The lack of accurate predictive models, which are essential for the design of high-performance laser-fusion experiments, is a major obstacle to achieving thermonuclear ignition. Here we report a statistical approach that was used to design and quantitatively predict the results of implosions of solid deuterium–tritium targets carried out with the 30-kilojoule OMEGA laser system, leading to tripling of the fusion yield to its highest value so far for direct-drive laser fusion. When scaled to the laser energies of the National Ignition Facility (1.9 megajoules), these targets are predicted to produce a fusion energy output of about 500 kilojoules—several times larger than the fusion yields currently achieved at that facility. This approach could guide the exploration of the vast parameter space of thermonuclear ignition conditions and enhance our understanding of laser-fusion physics.

Inertial confinement fusion (ICF)[1–3] uses drivers such as lasers or pulsed power to implode a millimetre-size capsule made of a layer of cryogenic deuterium (D) and tritium (T) enclosed within an ablator shell composed of low-atomic-number material (typically a CH polymer, high-density carbon or beryllium). The implosion produces copious amounts of DT fusion reactions, each producing a 14.1-MeV neutron and a 3.5-MeV α particle. Laser ICF uses either direct illumination of the capsule (direct drive)[4] or indirect illumination of the capsule via X-rays produced by laser heating of a high-atomic-number metal enclosure called a hohlraum (indirect drive)[5]. In direct-drive laser fusion, laser light directly incident on the capsule surface at intensities of $I \approx 10^{15}$ W cm^{-2} drives mass ablation, which accelerates the shell inwards at velocities from 300 to 500 km s^{-1}. As the shell converges, the central pressure increases and slows down the imploding capsule. The central core reaches temperatures of several kiloelectronvolts, and the heat flux from the centre ablates fuel off the inner shell surface to form a low-density (30–100 g cm^{-3}) central hot spot. This hot spot is confined by a nearly Fermi-degenerate[6] (200–1,000 g cm^{-3}) shell over a subnanosecond confinement time during which fusion reactions occur. If the areal density (ρR) of the compressed core is high enough, the α particles deposit their energy within the hot core, driving a feedback process (known as alpha heating)[7–9] that amplifies the fusion reaction rates. Under the appropriate conditions of pressure, temperature and confinement time[10–12], the alpha-heating process becomes intense enough to 'ignite' the central hot spot, resulting in a thermonuclear burn wave propagating radially through the entire dense fuel, producing fusion energy yields many times greater than the driver energy input. Until now, demonstrating ignition via laser fusion has been an elusive goal of the United States' ICF Program[13], although recent experiments[14,15] using indirect drive at the National Ignition Facility (NIF)[16] are approaching burning-plasma conditions[7,8],

where energy deposition by α particles is the dominant heating mechanism in the fusing plasma. Obtaining a burning plasma is considered to be an important intermediate milestone on the path to ignition. Recent publications provide detailed summaries of the indirect-[17,18] and direct-drive[19] efforts to achieve ignition, as well as an overview of the current status of ICF[13].

Historically, the design of ICF implosions has been primarily driven by one-dimensional (1D) and two-dimensional radiation–hydrodynamic (RH) codes[20–24]. Because of the substantial computational burden imposed by well resolved three-dimensional (3D) simulations, the use of 3D codes is mostly limited to post-shot simulations[25,26]. These simulations often include degradation sources (such as the ablator–fuel mix), engineering features (the stalk, fill tube or tent holding the capsule) and offsets from the target chamber centre to match experimental observables post hoc, but cannot yet reliably predict the results of a future experiment when changes to the target specifications and laser pulse shapes are made. At the Laboratory for Laser Energetics, the 1D RH code *LILAC*[23] is used to design direct-drive implosions, with experimentally benchmarked physics models for cross-beam energy transfer[27,28], thermal transport[28,29] and first-principles models of the equation of state[30]. In part because of these improvements in computer modelling, experiments carried out with the OMEGA laser[31] at energies of 26–28 kJ have demonstrated core compression that would produce substantial fusion α-particle self-heating when scaled hydrodynamically to NIF energies (1.9 MJ)[32,33]. Considerable performance improvements have also been achieved in indirect-drive ICF, with recent implosions of DT-layered targets producing over 50 kJ of fusion yield (about 2×10^{16} fusion reactions)[14], greatly exceeding the energy input to the fusion fuel.

Despite the remarkable progress in modelling, it is still not possible to predict a priori the results of an ICF experiment with enough accuracy

[1]Laboratory for Laser Energetics, University of Rochester, Rochester, NY, USA. [2]Department of Mechanical Engineering, University of Rochester, Rochester, NY, USA. [3]Department of Physics and Astronomy, University of Rochester, Rochester, NY, USA. [4]Dipartimento di Energetica, Politecnico di Milano, Milan, Italy. [5]Massachusetts Institute of Technology, Cambridge, MA, USA. *e-mail: vgop@lle.rochester.edu
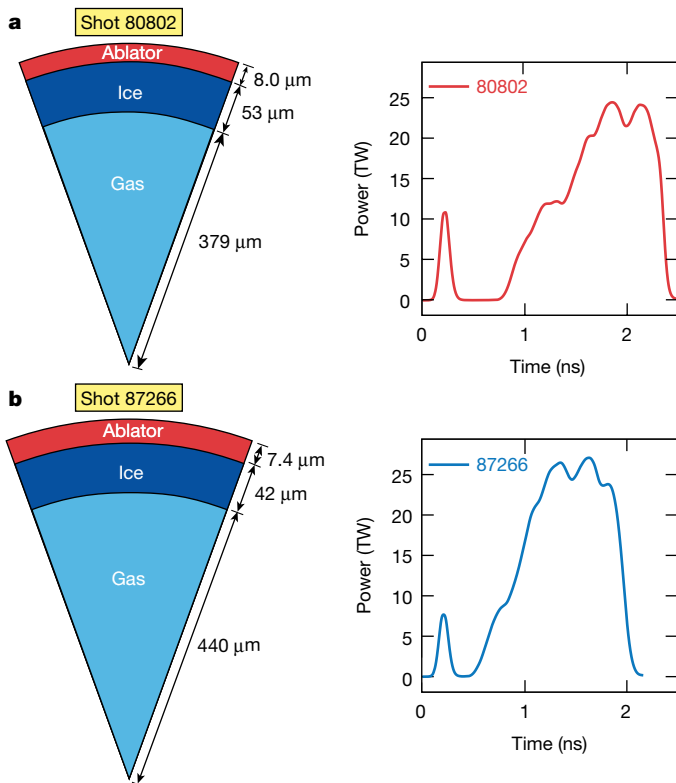
Fig. 1 | Target and laser pulse shape in the first phase of the Optimization Campaign. a, b, Shots 80802 (**a**) and 87266 (**b**) were the first and last shot, respectively, of the first phase of the Optimization Campaign. The implosion design was guided by statistical methods and led to the highest yield for direct-drive ICF. Design changes included larger target radii, reductions of the DT ice thickness (reflected in the changes to the gas, ice and ablator layer thicknesses shown on the left) and substantial changes to the laser pulse shape. The highest yield achieved in this campaign was three times higher than the previously held record[32].

to enable the implementation of iterative design methodologies and thus improve the implosion performance with a high degree of confidence. The lack of accurate prediction capability is a major obstacle in the quest for thermonuclear ignition. This paper describes a successful attempt to turn inaccurate code outputs into accurate predictions by finding statistical relationships between these outputs and experimental data. The application of this method has increased the fusion yield and performance in a two-phase, focused OMEGA experimental campaign (the Optimization Campaign; Fig. 1), leading to the highest fusion yield (about 1.6×10^{14} fusion reactions) so far in cryogenic DT direct-drive implosions (about three times higher than the previous yield record at the facility)[32], while achieving an areal density of about 160 mg cm^{-2}. Because the OMEGA laser lacks the energy to reach ignition, it is necessary to infer the performance of these implosions scaled to the energies that are currently accessible on the NIF. Such scaled targets are predicted to produce about 500 kJ of fusion yield (about 2×10^{17} fusion reactions) at 1.9 MJ of laser energy and symmetric illumination—several times higher than current high-performance implosions on the NIF[14], with further improvements expected to lead to megajoule yields at NIF energies.

Statistically driven implosion design

When designed in 1D, an ICF implosion is entirely specified by the target specifications and laser pulse shape (Fig. 1). The target specifications can be varied by changing the dimensions (target radius, ablator thickness and DT ice thickness) and the material properties (ablator compositions and DT vapour density). The laser pulse shape can be varied by changing the initial low-power portion, which is designed to launch carefully timed shock waves that set the entropy (or adiabat) of

the DT ice[34], or the ramp of the laser to full power, which is designed to adiabatically compress the shell. The late portion of the laser pulse can also be varied to prevent coasting of the shell or to launch a late shock as in shock ignition[35].

The unexplored parameter space of the laser pulse shape and target specifications is vast, and today's RH codes are not accurate enough to enable a high-confidence design of the optimum implosion that can be applied at the existing ICF laser facilities. In addition, the parameter space can be further extended if minor upgrades to the facility are included, such as changing the laser spot size on the target, the target mount or the DT filling procedure[19]. It is not clear at this time whether the lack of predictive capability in the RH codes is a result of inaccuracies in the physics models or of incomplete inputs provided to the codes or, as is most likely, both. It is clear, however, that the lack of predictive capability has hindered progress in improving the fusion yields in ICF. It is also clear that in this vast and unexplored parameter space, there almost certainly exists a combination of target specifications, laser pulse shape and minor upgrades that leads to considerably better performance than today's implosions.

Here, we describe a new, statistically driven framework to aid in the traversal of this space, predict and design high-performance implosions and speed up the quest for thermonuclear ignition. The fundamental principle behind this new framework is that even though the RH codes are inaccurate, the experimental observables $\boldsymbol{O}^{\mathrm{exp}}$ are expected to be correlated to the code output variables $\boldsymbol{O}^{\mathrm{sim}}$ because both the experiment and the code use the same input—that is, the laser pulse shape and the target geometry. Because the codes do not accurately reproduce the experimental results, the code–experiment relations are no longer one to one for each variable ($O_j^{\mathrm{exp}} \neq O_j^{\mathrm{sim}}$); instead, a more global relation is theorized to exist, for which each experimental observable is correlated to a combination of code output variables.

To elucidate this approach, we use a set of parameters $\boldsymbol{I}_{1\mathrm{D}}$ to denote the initial conditions determining the 1D implosion dynamics (that is, the laser pulse shape and target specifications) and $\boldsymbol{S}_{3\mathrm{D}}$ for the 3D non-uniformity seeds. The latter can be divided into two types: (1) 'random' initial seeds ($\boldsymbol{S}_{3\mathrm{D}}^{\mathrm{ran}}$) that vary from shot to shot because of circumstances outside the scientist's control (for example, roughness of the target surface and DT ice, power imbalance between laser beams, beam mispointing, beam mistiming and target offset from the target chamber centre), and (2) systematic seeds ($\boldsymbol{S}_{3\mathrm{D}}^{\mathrm{sys}}$) that are approximately invariant in every implosion (for example, laser speckle spectrum patterns, laser beam pattern on target, and perturbations from the target mount) because of the intrinsic nature of the facility. It can be useful to think of the random seeds as a measure of the reproducibility of experiments performed at the facility. The vast majority of experiments carried out at OMEGA are reproducible and exhibit negligible random seeds. Both the experiment and RH simulations can be thought of as highly nonlinear black boxes described by the sets of functions $\boldsymbol{F}_{\mathrm{exp}}$ and $\boldsymbol{F}_{\mathrm{sim}}$, respectively, so that

$$\boldsymbol{O}^{\mathrm{exp}} = \boldsymbol{F}_{\mathrm{exp}}\left[\boldsymbol{I}_{1\mathrm{D}}, \boldsymbol{S}_{3\mathrm{D}}^{\mathrm{sys}}, \boldsymbol{S}_{3\mathrm{D}}^{\mathrm{ran}}\right] \qquad (1)$$

$$\boldsymbol{O}_{1\mathrm{D}}^{\mathrm{sim}} = \boldsymbol{F}_{\mathrm{sim}}\left[\boldsymbol{I}_{1\mathrm{D}}, 0, 0\right] \qquad (2)$$

The zeros on the right-hand side of equation (2) indicate that only 1D RH codes are used in this analysis. Physically, the 3D nonuniformities are coupled into $\boldsymbol{F}_{\mathrm{exp}}$ and $\boldsymbol{F}_{\mathrm{sim}}$ via implosion hydrodynamic instabilities (Rayleigh–Taylor[36] and Richtmyer–Meshkov[37]). The growth of the instabilities is governed by the overall 1D hydrodynamic conditions, which are set by the 1D input conditions $\boldsymbol{I}_{1\mathrm{D}}$ representing the target specifications and the laser pulse shape.

Although it would be desirable to elucidate the form of $\boldsymbol{F}_{\mathrm{exp}}$ (either analytically or statistically), the scarcity of experimental data and the highly nonlinear nature of the physical processes that are probably involved make this unfeasible. A general expression for $\boldsymbol{F}_{\mathrm{sim}}$ is also difficult to find in an analytical form but it can be constructed via a
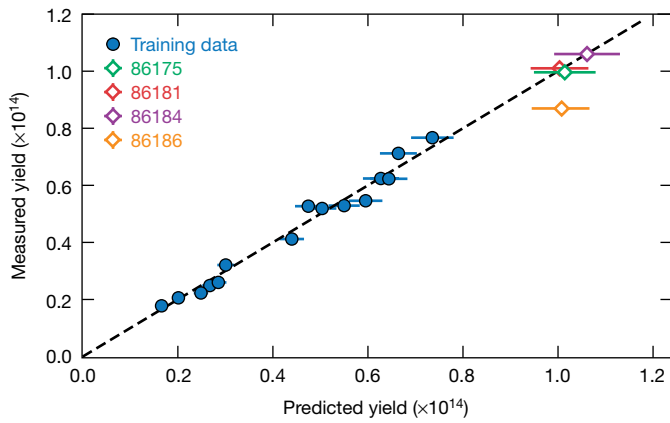
Fig. 2 | Prediction performance of the statistical model. The application of equation (6) to the experimental design yielded four high-yield implosions in July 2017 (open diamonds). The yield of three of these implosions was well predicted to within 5% of the corresponding measured values, whereas the yield of the fourth implosion (orange) was overpredicted because of a large target offset. Horizontal error bars and centre values represent the standard deviation and mean, respectively, of predictions made using 500 draws from the posterior distribution for the regression parameters.

large ensemble of numerical simulations[38]. However, if we assume that O_{1D}^{sim} is uniquely specified by a choice of I_{1D} through F_{sim}, then we can suppose that equation (2) can be inverted, leading to $I_{1D} = F_{sim}^{-1}(O_{1D}^{sim})$. Substituting into equation (1) leads to the relation

$$O^{exp} = F_{exp}[F_{sim}^{-1}(O_{1D}^{sim}), S_{3D}^{sys}, \ S_{3D}^{ran}] \qquad (3)$$

which can be considered as a mapping of the experimental observables onto the simulation outputs and the nonuniformity seeds. If experiments are repeatable (that is, the initial conditions that drive S_{3D}^{ran} are generally well controlled), the effect of random nonuniformities can be neglected with respect to the systematic ones, and equation (3) becomes a useful relation because S_{3D}^{sys} are constants for all experiments with a given set of systematic nonuniformities, leading to

$$O^{exp} \approx F_{exp}[F_{sim}^{-1}(O_{1D}^{sim}), S_{3D}^{sys}, \ 0] \qquad (4)$$

We note that in the limit where $S_{3D}^{sys} = 0$ and the RH codes perfectly model physical reality, equation (4) is reduced to $O^{exp} \approx O_{1D}^{sim}$. Although this is not the case with current RH codes, it is reasonable[27–30] to suppose that RH code outputs are tightly connected to the experimental observables so that the complicated relation in equation (4) can be approximated with simple functions, such as the power laws

$$O_j^{exp} \approx C_j(S_{3D}^{sys}) \prod_{i=1}^{N} [(O_{1D}^{sim})_i]^{\mu_i^j(S_{3D}^{sys})} \qquad (5)$$

where C_j and μ_i^j are constants for a given set of systematic nonuniformity seeds S_{3D}^{sys} and are determined by regressing against the existing experimental data. If the models in the RH codes are very inaccurate, equation (4) will need to be represented by more-complex basis functions or regression models. To design experiments using these statistical models, hydrodynamic simulations can be run for initial conditions constrained by considerations such as available resources, facility safety thresholds and physical intuition. Within these constraints, the space of initial conditions can then be sampled to optimize a metric predicted by the statistical models that are applied to these hydrodynamic simulations.

Experimental results

The set of implosion experiments designed in the first phase of the Optimization Campaign to test the statistical mapping are less

convergent than those of ref. [32], which until recently were considered to have the best performance, with fusion yields of about 4×10^{13} to 5×10^{13}. These less-convergent implosions were designed with the goal of steadily increasing the fusion yield. The simulated convergence ratio (the ratio between the outer radius of the target and the hot-spot radius) for this new implosion ensemble varies between 14 and 18, whereas the implosions of ref. [32] have a simulated convergence ratio of about 20–22. To illustrate the operation of the statistical model, we describe its application to the design of implosions carried out in July 2017 using experimental and simulated data from 15 implosions performed between October 2016 and April 2017. The ensemble of implosion experiments included plastic shells with an outer diameter of 870 μm, wall thickness of 7.5–8 μm and DT ice thickness varying from 42 to 53 μm. Each implosion was simulated with the 1D RH code *LILAC*. Bayesian inference[39] was used to find the most likely values for the coefficients C_j and μ_i^j in equation (5) for copper activation measurements of the neutron yield, giving

$$Y_{pred}^{exp} = \frac{4.2 \times 10^{13}}{\delta_{SSD}^{0.3}} \left(\frac{V_{sim}}{400}\right)^{4.2} \left(\frac{M_{sim}}{0.01}\right)^{0.6} \qquad (6)$$

where V_{sim} is the simulated implosion velocity in kilometres per second, M_{sim} is the simulated total stagnating mass at peak neutron rate in milligrams and δ_{SSD} is a categorical parameter denoting the presence ($\delta_{SSD} = 1$) or absence ($\delta_{SSD} = 2$) of smoothing by spectral dispersion (SSD)[40]. Posterior distributions for the regression parameters are shown in Extended Data Fig. 1. The choice of V_{sim} and M_{sim} as regression parameters is physically reasonable because the yield is a strong function of the internal energy of the hot spot and the inertia of the confining shell, which in turn are related to the kinetic energy of the imploding target and the shell mass, respectively. Strictly speaking (see ref. [10]), the simulated areal density should have been included in the formulation of equation (6), but because the simulated areal density did not vary considerably (with a coefficient of variation of 8%) within the ensemble of data, its inclusion could not be justified. The use of δ_{SSD} is necessary because SSD reduces the short-wavelength seeds for hydrodynamic instabilities (which is often referred to as laser imprinting[41]) and therefore changes the systematic seeds of nonuniformities. For simplicity, we assume that SSD only changes the constant C_j in equation (5). Figure 2 compares the measured fusion yield of the October 2016–April 2017 shots (blue circles) with the predicted values from equation (6) for both the training data and for the July experiments.

Equation (6) indicates that an increase in either simulated stagnated mass or implosion velocity would lead to higher fusion yields. This would be achievable by increasing the outer diameter of the target to increase the laser–target coupling while keeping the ice thickness constant. Accordingly, the July 2017 experiments used shells with outer diameter of 910 μm and ice thickness of 42 μm driven by 28 kJ of ultraviolet light. The fusion yields from these shots are shown as diamonds in Fig. 2. Three out of the four shots fell within 5% of the predictions from equation (6). The fourth shot was 15% below the predicted value because of an unexpected large initial target offset (38 μm, compared to typical values of about 5 μm). Large target offsets are an example of rare random events that result from vibrations of the target holder (a long stalk glued to the plastic shell). Equation (6) assumes negligible values of S_{3D}^{ran}; therefore, it cannot account for the effects of these events unless appropriately modified.

On the basis of the successful application of equation (6) to the experiments of July 2017, more implosions were included into the statistical model. The expanded dataset was not limited to the low-convergence implosions used to formulate equation (6) and spanned a wider range of initial conditions with simulated convergence ratios varying from 14 to 22, including the highest-convergence implosions[32]. With the expanded dataset, it became possible to infer the dependence of the yield on a larger set of code output parameters, updating the inference as additional data were collected. This resulted in the following model for the measured yield
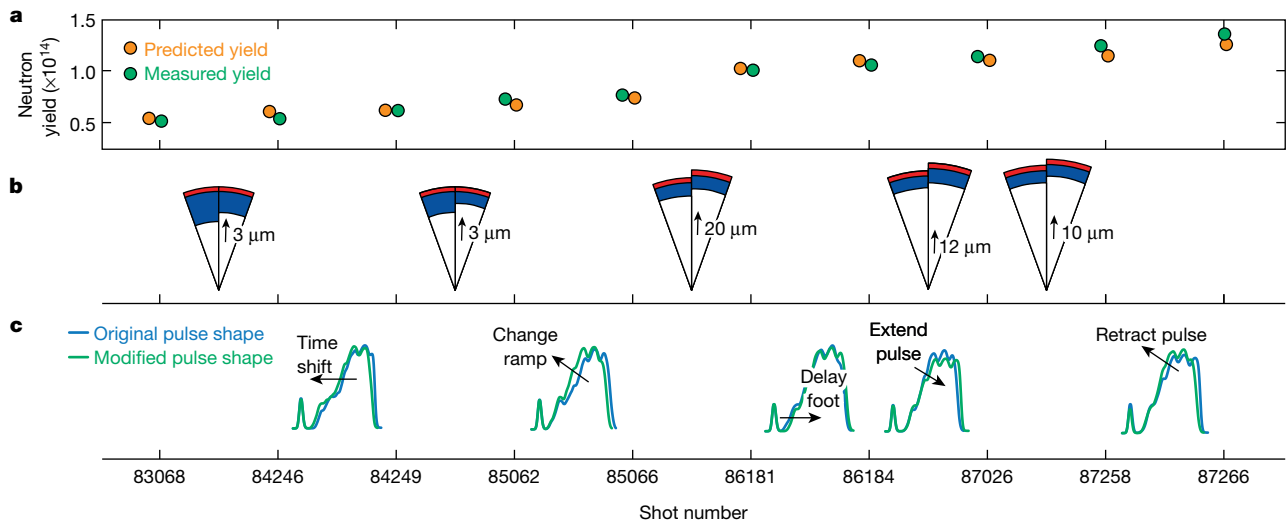
Fig. 3 | Application of the statistical model to the experimental design of the DT target. a–c, Prediction performance (**a**), target changes (**b**) and pulse shape changes (**c**) for implosions performed in October 2016–October 2017. As shown in **a**, the measured neutron yield (green circles) is accurately predicted by the statistical model (orange circles). In **b**, the gas, ice and ablator layers are shown in white, blue and red, respectively. Pulse shape modifications are shown in **c** by the blue (original) and green (modified) curves. The ice was thinned to reduce the total mass of the target, and the outer radius was increased to enhance the coupled laser energy, allowing the shell to reach a higher velocity and yield, in agreement with equation (7). The main pulse was modified to maintain the shock timing and reduce the time between the end of the laser pulse and stagnation. Modifications to the pulse shape in **c** are highlighted, with the arrows pointing in the direction of change of the pulse shape. Major modifications to the target dimensions are highlighted in **b**, with arrows pointing in the direction of change in ice thickness (83068 to 84246 and 84249 to 85062) or outer radius (subsequent shots).

$$Y_{\mathrm{pred}}^{\mathrm{exp}} = \frac{4.4 \times 10^{13}}{\delta_{\mathrm{SSD}}^{0.35} R_T^{0.6}} \left(\frac{V_{\mathrm{sim}}}{400} \right)^4 \left(\frac{M_{\mathrm{sim}}}{0.01} \right)^{0.6} \left(\frac{\rho R_{\mathrm{sim}}}{110} \right)^{0.3} \left(\frac{415}{R_0} \right)^{0.6} \quad (7)$$

where ρR_{sim} is the simulated total areal density at stagnation in milligrams per square centimetre, R_0 is the initial outer radius of the target in micrometres and $R_T = T_{\max}/T_{\min}$ represents the variation in the measured apparent ion temperatures along different lines of sight (typically an effect of about 10% on the yield). Posterior distributions for the regression parameters are shown in Extended Data Fig. 2. Because the expanded training data had a reasonable variation in the simulated areal density and outer radius, they were used when formulating the model. Differences in apparent ion-temperature measurements in the neutron time-of-flight detectors[42] are in part due to residual bulk motion or flows inside the central hot spot[43]; such flows are caused by large, low-ℓ mode asymmetries[44] because of random events, such as target offsets, which are an issue for a small but nontrivial number of implosions. The parameter R_T is therefore a measure of these random low-ℓ-mode asymmetries because it acts as a proxy for the effects of non-zero $\mathbf{S}_{3D}^{\mathrm{ran}}$ in equations (3) and (4). Although its inclusion is not strictly necessary to formulate equation (7), it allows us to increase the size of the training and validation data that are used to formulate and validate the model. A baseline value of $R_T = 1.15$ observed for nominal implosions, corresponding to a yield degradation of roughly 10% (see Extended Data Fig. 3), is used when making predictions with such models, and can be thought of as a prior on the facility's ability to constrain the initial conditions that drive R_T.

The accuracy of equation (7) in predicting the fusion yields from some representative changes in target specifications and laser pulse shape during the first phase of the Optimization Campaign is shown in Fig. 3. For each change of pulse shape and target size, the statistical model accurately predicted the measured yield (orange and green circles in Fig. 3a). The accuracy of the model against all of the data (in blue), including those affected by random nonuniformities, compared to the simulated 1D yield from *LILAC* (in red) is shown in Fig. 4. The large discrepancy between the measured and *LILAC*-simulated yields shown in Fig. 4 highlights the ability of the statistical model to turn inaccurate code results into accurate predictive tools. Using 1D RH simulations, combined with a constantly varying version of equation (7), progressively higher yield implosions were designed and carried out at OMEGA to achieve a yield of about 1.4×10^{14} fusion reactions with an areal density of 100 mg cm^{-2} for shot 87266 (Fig. 1b).

During the second phase of the Optimization Campaign, the framework was applied to increase the areal density in concert with the yield. Because the areal density, unlike the yield, depends strongly on the details of the pulse shape, this process was more challenging and required that the design be driven exclusively by the results of the statistical model. Once again using 1D RH simulations, along with updated versions of equation (7) and an analogous expression for the areal density (Extended Data Figs. 4, 5), implosions were designed and carried out at OMEGA. The best-performing shot of this phase, 90288, increased the yield slightly to about 1.6×10^{14} and the areal density by 60% to 160 mg cm^{-2} compared with shot 87266. Not only is this the highest-yield, direct-drive cryogenic implosion, but it also exhibits the highest normalized Lawson triple product for direct drive[8]. A detailed description of these high-performance designs will be presented in a future publication.

Extrapolated performance and outlook

The performance of OMEGA implosions is typically assessed by extrapolating their fusion yield to NIF laser energies and powers. Hydrodynamic scaling relations[33,45] are used to estimate the fusion yield at the NIF energy scale if the same core conditions of hot-spot pressure and shell density achieved with OMEGA are reproduced at the NIF within a larger volume and mass from implosions that match NIF laser energy and power. The effects of alpha heating, which are negligible at lower driver energies but important at the NIF energy scale, are also considered. Hydrodynamic scaling to NIF energies can be carried out using simple analytical estimates[45] or detailed computer simulations[33], both of which lead to similar results because of the simplicity and robustness of the hydrodynamic scaling. Following ref. [33], two-dimensional simulations were used to reproduce the same core conditions observed in the OMEGA experiments, starting from 1D simulations that were subsequently degraded by applying nonuniformity spectra constructed to reproduce the observed core distortions, hot-spot volume, X-ray emission profile, ion temperature, burnwidth,
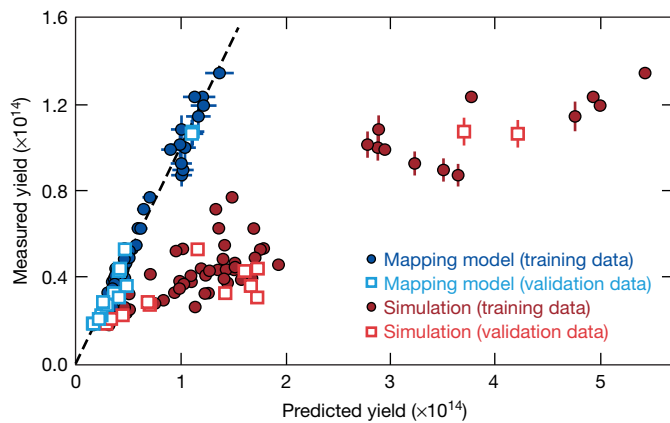
Fig. 4 | Comparison of prediction accuracies of the statistical model and the simulated 1D yield. The fusion yield from the statistical model (blue) in equation (7) accurately predicts the experimental yields for both low- and high-yield implosions from the first phase of the Optimization Campaign in both the training (circles) and validation (squares) data. By comparison, the simulated 1D yields from *LILAC* (red) show very large errors and overpredict the potential gains in yields. The horizontal error bars represent one standard deviation for predictions from the posterior distribution. Vertical error bars represent one standard deviation for the neutron yield detector. The dashed line is a guide for the eye.

areal density, fusion yield and inferred pressure. Such simulated core conditions approximately matched all experimental observables. These simulated implosions were then scaled to match the laser energy and power of the NIF and used to benchmark the analytical formulas[33]. Although it may be possible for the same core conditions to be reproduced by a different choice of nonuniformities or different 1D dynamics[46], hydrodynamic scaling leads to similar extrapolated fusion yields because the fusion yield depends primarily on the stagnation properties of the implosion and is insensitive to how such properties are achieved.

The process of hydrodynamic scaling assumes the same laser-to-capsule energy coupling and the same laser–plasma interaction physics[47]. It does not account for the differences in laser beam geometry between OMEGA and the NIF. OMEGA uses a symmetric illumination pattern suitable for direct-drive ICF whereas the NIF currently uses a polar configuration that is more suitable for indirect drive. It follows that the extrapolated yields should be considered as an upper bound and must be reduced to account for the lack of illumination symmetry at the NIF (unless the NIF is converted to symmetric drive and the changes in laser-to-capsule coupling and laser–plasma interaction physics are properly understood at this scale). There is an extensive research effort underway at OMEGA and the NIF to explore and quantify these effects.

The fusion yields from the Optimization Campaign, hydrodynamically scaled to 1.9 MJ of symmetric laser illumination, would produce about 500 kJ of fusion yield, as shown in Fig. 5. By comparison, the previous best-performing implosions of ref. [32] hydrodynamically scaled to 1.9 MJ would produce a fusion yield of about 100 kJ (Fig. 5). When compared to current indirect-drive implosions, which produce up to 56 kJ of fusion energy[14], the fusion yield from hydrodynamically scaled OMEGA implosions is substantially higher, mostly because of the larger size of the fusion core, with alpha heating causing yield amplification by a factor of about 3 in both cases. As shown in Fig. 5, a modest increase in areal density or yield (blue arrow) will result in OMEGA core conditions that scale to megajoule fusion yields at NIF laser energies. The next step of the direct-drive programme at the OMEGA laser is therefore to apply the statistical model described in this paper to further increase the areal density by increasing the convergence of the target while keeping the yields at their highest levels. Given the vast unexplored parameter space of laser pulse shapes, target specifications and facility upgrades, it is very likely that optimized OMEGA implosions will outperform the encouraging results reported
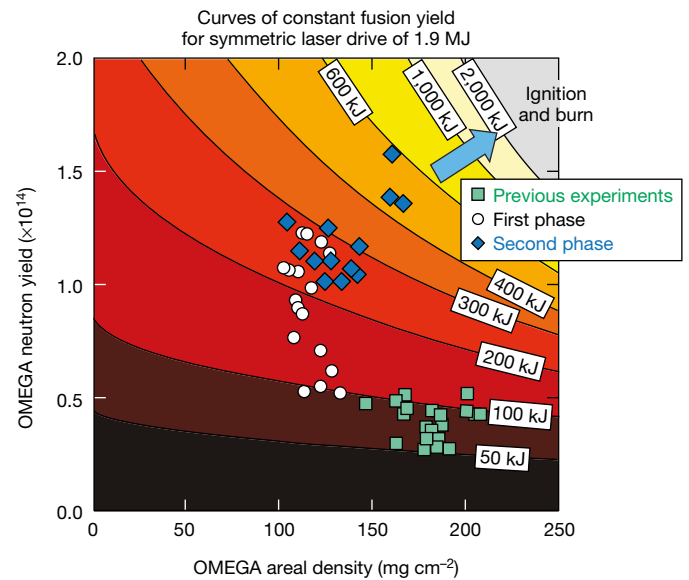
Fig. 5 | NIF-extrapolated fusion yields for selected OMEGA implosions. The results of the first (white circles) and second (blue diamonds) phases of the Optimization Campaign are compared with those of previous experiments (green squares)[32]. The black lines represent curves of constant extrapolated fusion yield for symmetric laser drive of 1.9 MJ. The blue arrow shows the direction of future experiments. The increase in areal density and yield achieved using the framework presented in this paper results in implosions that would produce about 2×10^{17} fusion reactions, or 500 kJ of fusion energy, at NIF energies of about 1.9 MJ.

in this paper. Although we cannot predict the final magnitude of the expected improvements in performance, we are confident that the new predictive capability provided by statistical mapping, as well as the high reproducibility and throughput of OMEGA, will provide an accurate and reliable framework that will enable rapid progress towards demonstrating thermonuclear ignition and burn in the laboratory.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, statements of data availability and associated accession codes are available at https://doi.org/10.1038/s41586-019-0877-0.

1. Nuckolls, J., Wood, L., Thiessen, A. & Zimmerman, G. Laser compression of matter to super-high densities: thermonuclear (CTR) applications. *Nature* **239**, 139–142 (1972).
2. Atzeni, S. & Meyer-ter-Vehn, J. *The Physics of Inertial Fusion: Beam Plasma Interaction, Hydrodynamics, Hot Dense Matter* 1st edn (Oxford Univ. Press, Oxford, 2004).
3. Slutz, S. A. et al. Pulsed-power-driven cylindrical liner implosions of laser preheated fuel magnetized with an axial field. *Phys. Plasmas* **17**, 056303 (2010).
4. Craxton, R. S. et al. Direct-drive inertial confinement fusion: a review. *Phys. Plasmas* **22**, 110501 (2015).
5. Lindl, J. D. *Inertial Confinement Fusion: The Quest for Ignition and Energy Gain Using Indirect Drive* (Springer-Verlag, New York, 1998).
6. Giorgini, S., Pitaevskii, L. P. & Stringari, S. Theory of ultracold atomic Fermi gases. *Rev. Mod. Phys.* **80**, 1215–1274 (2008).
7. Christopherson, A. R. et al. A comprehensive alpha-heating model for inertial confinement fusion. *Phys. Plasmas* **25**, 012703 (2018).
8. Betti, R. et al. Alpha heating and burning plasmas in inertial confinement fusion. *Phys. Rev. Lett.* **114**, 255003 (2015).
9. Hurricane, O. A. et al. Inertially confined fusion plasmas dominated by alpha-particle self-heating. *Nat. Phys.* **12**, 800–806 (2016).
10. Betti, R. et al. Thermonuclear ignition in inertial confinement fusion and comparison with magnetic confinement. *Phys. Plasmas* **17**, 058102 (2010).
11. Chang, P. Y. et al. Generalized measurable ignition criterion for inertial confinement fusion. *Phys. Rev. Lett.* **104**, 135002 (2010).
12. Spears, B. K. et al. Performance metrics for inertial confinement fusion implosions: aspects of the technical framework for measuring progress in the National Ignition Campaign. *Phys. Plasmas* **19**, 056316 (2012).

13. Betti, R. & Hurricane, O. A. Inertial-confinement fusion with lasers. *Nat. Phys.* **12**, 435–448 (2016); corrigendum **12**, 717 (2016).
14. Le Pape, S. et al. Fusion energy output greater than the kinetic energy of an imploding shell at the National Ignition Facility. *Phys. Rev. Lett.* **120**, 245003 (2018).
15. Hurricane, O. A. et al. Fuel gain exceeding unity in an inertially confined fusion implosion. *Nature* **506**, 343–348 (2014); corrigendum **510**, 432 (2014).
16. Campbell, E. M. & Hogan, W. J. The National Ignition Facility – applications for inertial fusion energy and high-energy-density science. *Plasma Phys. Contr. Fusion* **41**, B39–B56 (1999).
17. Edwards, M. J. et al. Progress towards ignition on the National Ignition Facility. *Phys. Plasmas* **20**, 070501 (2013).
18. Lindl, J. et al. Review of the National Ignition Campaign 2009–2012 *Phys. Plasmas* **21**, 020501 (2014); erratum **21**, 129902 (2014).
19. Goncharov, V. N. et al. National direct-drive program on OMEGA and the National Ignition Facility. *Plasma Phys. Contr. Fusion* **59**, 014008 (2017).
20. Marinak, M. M. et al. Three-dimensional HYDRA simulations of National Ignition Facility targets. *Phys. Plasmas* **8**, 2275–2280 (2001).
21. Zimmerman, G., Kershaw, D., Bailey, D. & Harte, J. The *LASNEX* code for inertial confinement fusion. *J. Opt. Soc. Am.* **68**, 549 (1978).
22. Radha, P. B. et al. Two-dimensional simulations of plastic-shell, direct-drive implosions on OMEGA. *Phys. Plasmas* **12**, 032702 (2005).
23. Delettrez, J., Epstein, R., Richardson, M. C., Jaanimagi, P. A. & Henke, B. L. Effect of laser illumination nonuniformity on the analysis of time-resolved X-ray measurements in UV spherical transport experiments. *Phys. Rev. A* **36**, 3926–3934 (1987).
24. Atzeni, S. et al. Fluid and kinetic simulation of inertial confinement fusion plasmas. *Comput. Phys. Commun.* **169**, 153–159 (2005).
25. Clark, D. S. et al. Radiation hydrodynamics modeling of the highest compression inertial confinement fusion ignition experiment from the National Ignition Campaign. *Phys. Plasmas* **22**, 022703 (2015).
26. Igumenshchev, I. V. et al. Three-dimensional modeling of direct-drive cryogenic implosions on OMEGA. *Phys. Plasmas* **23**, 052702 (2016).
27. Igumenshchev, I. V. et al. Crossed-beam energy transfer in implosion experiments OMEGA. *Phys. Plasmas* **17**, 122708 (2010).
28. Michel, D. T. et al. Measurements of the ablation-front trajectory and low-mode nonuniformity in direct-drive implosions using X-ray self-emission shadowgraphy. *High Power Laser Sci. Eng.* **3**, e19 (2015).
29. Goncharov, V. N. et al. Early stage of implosion in inertial confinement fusion: shock timing and perturbation evolution. *Phys. Plasmas* **13**, 012702 (2006).
30. Hu, S. X. et al. First-principles equation of state of polystyrene and its effect on inertial confinement fusion implosions. *Phys. Rev. E* **92**, 043104 (2015).
31. Boehly, T. R. et al. Initial performance results of the OMEGA Laser System. *Opt. Commun.* **133**, 495–506 (1997).
32. Regan, S. P. et al. Demonstration of fuel hot-spot pressure in excess of 50 Gbar for direct-drive, layered deuterium–tritium implosions on OMEGA. *Phys. Rev. Lett.* **117**, 025001 (2016).
33. Bose, A. et al. Core conditions for alpha heating attained in direct-drive inertial confinement fusion. *Phys. Rev. E* **94**, 011201 (2016).
34. Robey, H. F. et al. Measurement of high-pressure shock waves in cryogenic deuterium-tritium ice layered capsule implosions on NIF. *Phys. Rev. Lett.* **111**, 065003 (2013).
35. Betti, R. et al. Shock ignition of thermonuclear fuel with high areal density. *Phys. Rev. Lett.* **98**, 155001 (2007).
36. Kilkenny, J. D. et al. A review of the ablative stabilization of the Rayleigh–Taylor instability in regimes relevant to inertial confinement fusion. *Phys. Plasmas* **1**, 1379–1389 (1994).
37. Brouillette, M. The Richtmyer–Meshkov instability. *Annu. Rev. Fluid Mech.* **34**, 445–468 (2002).
38. Peterson, J. L. et al. Enhancing hohlraum design with artificial neural networks. *Bull. Am. Phys. Soc.* **62** http://meetings.aps.org/link/BAPS.2017.DPP.BO7.6 (2017).
39. Salvatier, J., Wiecki, T. V. & Fonnesbeck, C. Probabilistic programming in python using PyMC₃. *PeerJ Comput. Sci.* **2**, e55 (2016).
40. Boehly, T. R. et al. Optical and plasma smoothing of laser imprinting in targets driven by lasers with SSD bandwidths up to 1 THz. *Phys. Plasmas* **8**, 2331–2337 (2001).
41. Hu, S. X. et al. Understanding the effects of laser imprint on plastic-target implosions on OMEGA. *Phys. Plasmas* **23**, 102701 (2016).
42. Murphy, T. J., Chrien, R. E. & Klare, K. A. Interpretation of neutron time-of-flight signals from current-mode detectors. *Rev. Sci. Instrum.* **68**, 610–613 (1997).
43. Murphy, T. J. The effect of turbulent kinetic energy on inferred ion temperature from neutron spectra. *Phys. Plasmas* **21**, 072701 (2014).
44. Woo, K. M. et al. Effects of residual kinetic energy on yield degradation and ion temperature asymmetries in inertial confinement fusion implosions. *Phys. Plasmas* **25**, 052704 (2018).
45. Nora, R. et al. Theory of hydro-equivalent ignition for inertial fusion and its applications to OMEGA and the National Ignition Facility. *Phys. Plasmas* **21**, 056316 (2014).
46. Bose, A. et al. Analysis of trends in implosion observables for direct-drive cryogenic implosions on OMEGA. *Phys. Plasmas* **25**, 062701 (2018).
47. Smalyuk, V. A. et al. Role of hot-electron preheating in the compression of direct-drive imploding targets with cryogenic D_2 ablators. *Phys. Rev. Lett.* **100**, 185005 (2008).

METHODS

Experimental facility. The implosions described here were performed at the University of Rochester's Omega Laser Facility[31] at the Laboratory for Laser Energetics (LLE). The implosions were driven by 60 symmetrically aligned ultra-violet laser beams with a wavelength of 350 nm and a spot radius of 415 μm, which were incident on a spherical plastic shell with a radius of 400–500 μm and an inner layer of DT ice of 41–53 μm thickness. The total laser energy delivered was 25–29 kJ, with a shaped pulse over 2–3 ns and peak power of 20–30 TW. Laser ablation drove the shell inwards until, at stagnation, a central hot spot at 3–5 keV was formed and was confined for 60–100 ps by the surrounding unablated DT plasma. Several important diagnostics were used to infer the initial and stagnation conditions of the implosion that were used in the statistical model and to determine the reasons for deviation of the measurements from the model:

1. Measurements of the total number of 14-MeV fusion neutrons were made with the neutron time-of-flight (nTOF)[42] and copper activation detectors and were used to infer the number of fusion reactions in the implosion.

2. Measurements of the temperature of the DT fusion reactants in the centre-of-mass reference frame were made with the nTOF detectors and were used to infer the ion temperature of the hot spot in the laboratory reference frame. Such measurements can be affected by bulk fluid motion of the plasma along the line of sight of the detector[43]. As a result, the variation of the ion temperature contains information about the asymmetry of the imploded target.

3. Measurements of the areal density of the confining shell were made by examining the backscatter of the 14-MeV neutrons off the shell in an nTOF detector[48] and the forward scatter of the 14-MeV neutrons off the shell in a magnetic recoil spectrometer[49].

4. Measurements of the initial target offset were made by observing the target with a high-speed video camera before it imploded and from the relative position of the X-ray self-emission from the compressed core.

5. Time-resolved measurements of the laser power in each of the beams were made using a set of multichannel ultraviolet streak cameras[50].

6. Shadowgraphy measurements[51] of the target provided the mean thicknesses and surface quality metrics for each layer.

Simulations. Simulations were carried out using the RH code *LILAC*, which was developed at the LLE, to study laser-driven implosions. *LILAC* is a 1D Lagrangian hydrocode with nonlocal thermal transport[29], multigroup radiation diffusion and first-principles equations of state[30,52]; it models laser deposition with 3D ray tracing, including a cross-beam energy transport model[27]. *LILAC* has been bench-marked extensively against experiments and has been shown to reproduce many important experimental quantities in several cases[28,30]. The initial conditions for each simulation were chosen to match the experiment as closely and systematically as possible in one dimension. The target layer thickness was the average value over the surface and the laser power was the average across all 60 beams. This ensured that the results from the simulations were internally consistent and reproduced the as-shot conditions as well as possible.

Statistical methods. A combined database with initial conditions, post-processed simulation data and experimental data for each shot was constructed for a large set of cryogenic DT implosions. Under the assumption that *LILAC* approximately reproduces the 1D dynamics and that the dominant 3D perturbations were systematic, it was shown that a predictive relationship should exist between *LILAC* post-processed and experimentally observed quantities. Power-law relations were chosen because of their simplicity and their ability to reproduce simulated performance metrics (yields and areal densities) using other simulated parameters[53]. Models were constructed using the PyMC3 Bayesian inference package[39] with a Gaussian likelihood function

$$L(Y, \mu, \varepsilon) \propto e^{-(Y-m)^2/(\varepsilon\mu)^2} \qquad (8)$$

where Y is the measured quantity, μ is the corresponding predicted value given a particular choice of model parameters and ε is a fractional error term that accounts for the heteroscedasticity of the measurement error. The priors for each of the model parameters were strong and normally distributed around their theoretical values, where available, and uninformative otherwise. Roughly 20% of the data were separated before inference into a validation set. The posterior distributions were sampled using the no-U-turn sampler[54] and were then used to make predictions on the validation set (as well as the data used for inference). The models were then evaluated on the basis of the mean absolute per cent error of the hold-out set, the width of the posterior distributions, the convergence and autocorrelation of each Markov chain Monte Carlo trace, and whether the 3σ credible interval for each parameter contained 0 (which led to immediate rejection). To make predictions and design experiments, *LILAC* simulations were run for initial conditions constrained in part by available resources and laser safety thresholds. Within these constraints, the pulse shape and targets were varied to optimize the yield predicted by the statistical models applied to *LILAC*.

Details of hydrodynamic scaling. When implosions carried out at OMEGA are extrapolated to the NIF, it becomes necessary to consider the effects of alpha heating, which are negligible at the lower driver energies of OMEGA. Hydrodynamic scaling to the NIF can be carried out through RH simulations or by using simple analytical models.

In the analytical model[53], hydrodynamic scaling requires implosions to be scaled up in size to match the equivalent NIF implosion so that it exhibits the same velocity and adiabat as the OMEGA implosion, with $M_{NIF} = \Delta^3 M_\Omega$, $R_{NIF} = \Delta R_\Omega$ and $\tau_{NIF} = \Delta \tau_\Omega$, where $\Delta = (E_{NIF}/E_\Omega)^{1/3}$ and M, R, τ and E are the target mass, length scale, time scale and driver energies, respectively, of the NIF and OMEGA (Ω) facilities. Under these conditions, the OMEGA and NIF implosions are deemed hydro-equivalent because all the hydrodynamic dimensionless parameters are identical. This leads to a Lawson triple product $2nT\tau = P\tau$ that scales with size (τ is the disassembly time, $\tau \approx R$, and P is the pressure, $P = nT$ for an ideal gas) and a yield $Y_{no\,\alpha}$ (without accounting for the effects of alpha heating) that scales as R^3 from the increased volume, as R from the longer burn time, and as $R^{0.3}$ from the increased temperature as a result of the larger volume-to-surface ratio; this leads to a total scaling of $R^{4.3}$. The effect of alpha heating[10] can be included by estimating the yield amplification from α particles through the factor $(1 - 1.04\Delta\chi_{no\alpha})^{-0.75}$, where $\chi_{no\alpha}$ is the normalized Lawson parameter for ICF, which can be expressed as $\chi_{no\alpha} = \rho R_{no\alpha}^{0.61}[1.2 \times 10^{-17}(Y_{no\alpha}/M_{stag})]^{0.34}$, with M_{stag}, $\rho R_{no\alpha}$ and $Y_{no\alpha}$ representing the stagnating mass in milligrams, the areal density in grams per square centimetre and the yield, neglecting alpha heating, respectively[10]. Combining the effects of alpha heating and size scaling leads to the total hydrodynamically scaled yield of

$$Y_{NIF} = \frac{Y_\Omega (E_{NIF}/E_\Omega)^{1.43}}{[1 - 1.04\chi_{no\alpha}^\Omega (E_{NIF}/E_\Omega)^{0.34}]^{0.75}} \qquad (9)$$

where Y_Ω is the yield measured in the OMEGA experiments, Y_{NIF} is the extrapolated yield to NIF laser energies, E_{NIF} and E_Ω are the laser energies of the NIF and OMEGA scale targets, and $\chi_{no\alpha}^\Omega$ is the normalized Lawson parameter of the OMEGA scale target.

In the simulation model, which is described in greater detail in ref. [33], the measured OMEGA core properties are reproduced using two-dimensional simulations with an ad hoc spectrum of nonuniformities. The strategy for the choice of perturbation is based on the observation that low-ℓ modes ($\ell < 6$) reduce the hot-spot pressure by increasing its volume whereas mid-ℓ modes ($6 < \ell < 40$) reduce the hot-spot volume without affecting the pressure. Therefore, if the experimentally inferred pressure is below the 1D simulated value, we use low-ℓ modes to degrade the implosions. If the use of low-ℓ modes leads to simulated hot-spot volumes larger than the measured ones, we add mid-ℓ modes to reduce the simulated volume. The simulations are scaled up in size to match the NIF energy as in the analytical model and they are repeated including α-particle deposition and transport. Laser imprinting can be approximated as 1D degradation in 1D codes. The imprinting front penetrates inwards from the ablation front, and its effect can be interpreted as an effective reduction of outer-shell mass[55].

The proposed hydrodynamic scaling is robust and rather insensitive to the type of degradation used as long as all experimental observables are reproduced in the reconstruction of the compressed core, and the analytical scaling relations have shown good agreement with detailed numerical simulations of hydrodynamically scaled implosions[33].

The validity of the hydrodynamic scaling requires the coupling of the laser energy to the target and any degradation mechanism caused by laser–plasma instabilities or other non-hydrodynamic physics to be the same for OMEGA and for the NIF. At this stage, it is not possible to provide conclusive evidence regarding the scaling of laser-to-target coupling and laser–plasma instabilities from OMEGA to the NIF. Because of the larger plasma scale length, collisional absorption is higher for the NIF, which would improve energy coupling. However, simulations indicate that cross-beam energy transfer (CBET) could be slightly enhanced at the NIF scale, reducing energy coupling, although recent developments on wavelength detuning at the NIF have demonstrated a reduction of CBET in direct-drive implosions[56]. Laser–plasma instabilities are dominated by the two-plasmon decay (TPD) instability at OMEGA, leading to substantial hot-electron preheating[47] when the laser intensity exceeds 10^{15} W cm^{-2}. However, initial direct-drive experiments at the NIF indicate that the stimulated Raman scattering process is dominant[57] and the suprathermal electron temperature is lower than TPD electrons. Because colder electrons are more easily stopped before reaching the inner fuel layer, preheating from electrons excited by stimulated Raman scattering is less than that from TPD-excited electrons for the same suprathermal electron energy. As a result of these conflicting issues, it is not possible at this time to provide quantitative predictions for the scaling of laser–plasma interactions from OMEGA to the NIF. Direct-drive experiments are currently underway at the NIF with the specific goal of providing a quantitative assessment of laser–plasma instabilities at the NIF scale.

Physical interpretation and applicability of the statistical model. It is tempting to interpret the results of equations (6) and (7) physically and to compare them to analytical scaling relations for direct-drive ICF[52]. Although we believe that the framework presented in this paper can potentially be used to gain a deeper understanding of the physical processes that drive ICF implosions, there are some major issues that prevent a straightforward application towards this end. First and foremost is the fact that the inputs to equations (6) and (7) are not in fact physical quantities—they are the outputs from a 1D simulation code and are not guaranteed to be equivalent to the corresponding 1D physical quantities. There are multiple experimental indications that the 1D hydrocode *LILAC* does not accurately reproduce the 1D behaviour of the implosions. Second, equations (6) and (7) are the result of correlating a 1D hydrocode with 3D experimental results. As a result, differences between equations (6) and (7) and theoretical scaling relations will be caused partially by 3D effects and partially because the inputs to equations (6) and (7) do not exactly correspond to their experimental analogues. It is not trivial to separate these effects and directly identify the drivers for a given difference between theory and statistical model, although future work may shed light on this issue.

Owing to the reasons stated above, it is also worth noting that equations (6) and (7) are only valid for a carbon–deuterium ablator, spherical direct-drive cryogenic DT targets used at the OMEGA Laser Facility after its last major upgrade in 2014, using the results from *LILAC* v2017.0.1, and are not expected to be valid for other targets, fusion schemes, hydrocodes or facilities. However, it is important to emphasize that the statistical framework described by equations (1)–(5) is generally valid and can be applied at any experimental facility, as long as F_{sim} remains deterministic and invertible and S_{3D}^{ran} remains small compared to S_{3D}^{sys} for that facility.

Choice of regression variables. In principle, regression variables can be selected without any physical insight, using common-feature selection methods. This is however a risky venture because it has a serious risk of overfitting, even when regularization methods such as cross validation are used.

Instead, it is desirable to use variables that are deemed to be the most physically relevant. The simplest qualitative physics arguments in support of using V_{sim}, M_{sim} or ρR_{sim} in equations (6) and (7) are: (a) the implosion velocity determines the temperature[58] ($T_i \approx V_{sim}$) and the fusion reactivity is a strong function of T_i, (b) the stagnating mass determines the size of the system and the confinement time[10] ($\tau_{sim} \approx \sqrt{M_{sim}}$), and (c) the areal density measures convergence (related to the final pressure) and confinement time.

At a more quantitative level, one can ask how the simulated yield depends on the simulated parameters to further verify the use of variables in the mapping of the experiments. The dependence of the simulated yield on simulated variables was derived in ref. [58] using T_{sim}, ρR_{sim} and M_{sim}. A similar formula can be obtained using V_{sim} instead of T_{sim}, as shown in Extended Data Fig. 6 for simulations with the *LILAC* code. This indicates that V_{sim}, M_{sim} and ρR_{sim} are good parameters for predicting the yield, although the experimental mapping will probably require different power indices and even different parameters to account for 3D effects and for deficiencies in the physics models. Equation (7), for instance, not only uses very different power indices than those shown above, but it also uses a new parameter (the target radius R_0) that does not appear above. R_0 is probably required to correct the simulations for inaccuracies in energy coupling (laser to target) when the target diameter is varied. In particular, CBET is strongly affected by the ratio of the laser beam to the target sizes. By changing R_0, that ratio is varied and inaccuracies in the CBET modelling can be compensated by a parametric dependence on R_0, as shown in equation (7).

Usage of the a posteriori measurement R_T in statistical models. The inclusion of the ratio of the maximum to the minimum measured ion temperature, R_T, in equation (7) may seem to invalidate the stated purpose of the framework presented in this paper because it is an experimental quantity that can only be determined after the experiment. However, this is not the case. The term R_T is used to act as a proxy for unexpectedly large random nonuniformity seeds that affect the experiment. A good example of this is the initial target offset: when the target is irradiated by

the laser, it is held in place by a thin (15 μm) stalk that is subject to some vibration. As a result, the target will be slightly off-centre when the laser fires, creating a nonuniform laser drive on the target surface. Because of the heavy electromagnetic damping of the stalk, the target is typically less than 10 μm off-centre, in which case the nonuniformity contribution is negligible. However, on rare occasions the target can be extremely off-centre, as was the case for shot 86186 shown in Fig. 2, which was nearly 40 μm off-centre. In this case, the resulting nonuniformity cannot be neglected in equation (4), which would require us to discard these experiments. Because experimental data are extremely expensive, we would prefer to use as much of the available data as possible. Because the ion temperature is raised by bulk fluid motion in the hot spot, R_T is a proxy for the effects of asymmetries in the hot spot, and its usage enables shots like 86186 (Fig. 2) to be included when formulating equation (7). For the majority of shots, the inclusion of R_T has little to no effect (see Extended Data Fig. 3) because most shots have small offsets ($\leq$10 μm) and good beam power balance (<3% to 4%). Approximately 70% of shots have R_T below 1.2, with a median of 1.15 (presumably this is caused by systematic nonuniformities). Whereas the framework can be used to formulate models without the use of regression variables such as R_T (the primary effect being simply an increase in the scatter of the data), its inclusion also allows the experimentalist to determine whether it is productive to repeat an experiment. For instance, if a shot occurs with a large target offset and has a large value of R_T, the experimenter can set R_T to its nominal value of about 1.15 and carry out a risk assessment of whether a repeat is likely to result in a substantial increase in performance. Finally, when designing experiments, the experimentalist can choose a value (or distribution) of R_T that they think is likely—formally, this can be thought of as a prior on the facility's ability to execute an experiment that meets specifications. We typically use $R_T = 1.1$–1.15 when designing experiments.

Data availability
Raw data were generated at the LLE's OMEGA Laser Facility. Derived data supporting the findings of this study are available from the corresponding author upon reasonable request.

48. Forrest, C. J. et al. High-resolution spectroscopy used to measure inertial confinement fusion neutron spectra on OMEGA. *Rev. Sci. Instrum.* **83**, 10D919 (2012).
49. Casey, D. T. et al. The magnetic recoil spectrometer for measurements of the absolute neutron spectrum at OMEGA and the NIF. *Rev. Sci. Instrum.* **84**, 043506 (2013).
50. Donaldson, W. R., Boni, R., Keck, R. L. & Jaanimagi, P. A. A self-calibrating, multichannel streak camera for inertial confinement fusion applications. *Rev. Sci. Instrum.* **73**, 2606–2615 (2002).
51. Edgell, D. H. et al. Three-dimensional characterization of cryogenic target ice layers using multiple shadowgraph views. *Fus. Sci. Technol.* **49**, 616–625 (2006).
52. Hu, S. X., Militzer, B., Goncharov, V. N. & Skupsky, S. First-principles equation-of-state table of deuterium for inertial confinement fusion applications. *Phys. Rev. B.* **84**, 224109 (2011).
53. Zhou, C. D. & Betti, R. A measurable Lawson criterion and hydro-equivalent curves for inertial confinement fusion. *Phys. Plasmas* **15**, 102707 (2008); correction 16, 079905 (2009).
54. Hoffman, M. D. & Gelman, A. The no-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* **15**, 1593–1623 (2014).
55. Goncharov, V. N. et al. Improving the hot-spot pressure and demonstrating ignition hydrodynamic equivalence in cryogenic deuterium–tritium implosions on OMEGA. *Phys. Plasmas* **21**, 056315 (2014).
56. Marozas, J. A. et al. First observation of cross-beam energy transfer mitigation for direct-drive inertial confinement fusion implosions using wavelength detuning at the National Ignition Facility. *Phys. Rev. Lett.* **120**, 085001 (2018).
57. Rosenberg, M. J. et al. Origins and scaling of hot-electron preheat in ignition-scale direct-drive inertial confinement fusion experiments. *Phys. Rev. Lett.* **120**, 055001 (2018).
58. Zhou, C. D. & Betti, R. Hydrodynamic relations for direct-drive fast-ignition and conventional inertial confinement fusion implosions. *Phys. Plasmas* **14**, 072703 (2007).
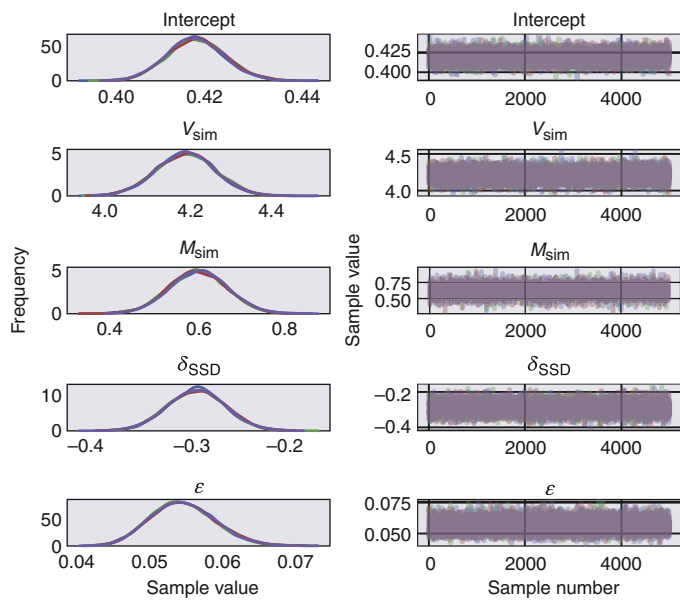
Extended Data Fig. 1 | Posterior predictive distributions for the parameters in equation (6). The intercept parameter refers to the constant C_j in equation (5), ε refers to the noise parameter in equation (8) and the remaining parameters are the μ_i values for the corresponding simulated quantities, as in equation (5). The Markov chains for each parameter are converged over 5,000 steps. The tails of the distributions have the same sign as the mean and are not excessively wide when compared to the mean. All these indicate that the model is well specified.
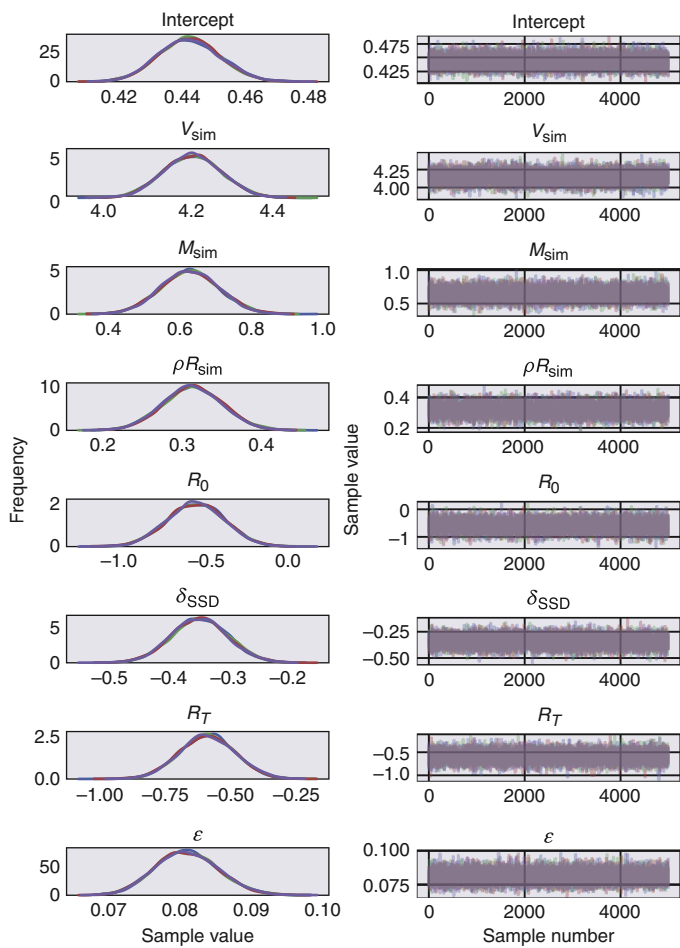
Extended Data Fig. 2 | Posterior predictive distributions for the parameters in equation (7). The intercept parameter refers to the constant C_j in equation (5), ε refers to the noise parameter in equation (8) and the remaining parameters are the μ_i values for the corresponding simulated quantities, as in equation (5). As in Extended Data Fig. 1, these results indicate that the model is well specified; we note that this does not guarantee that the model will have any predictive power.
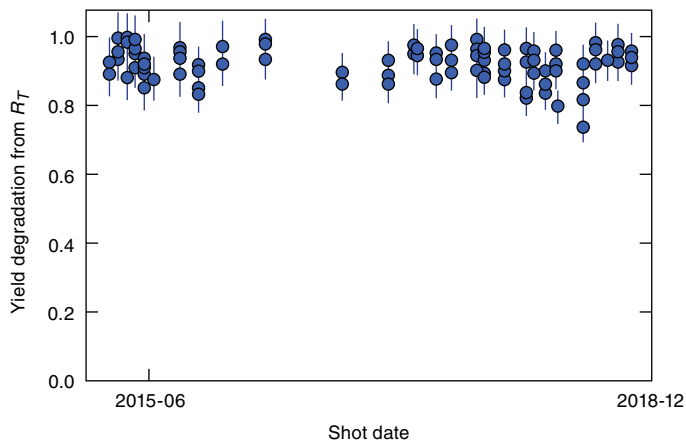
Extended Data Fig. 3 | The effect of R_T on the yield is less than 10% for the vast majority of implosions. Vertical error bars are one standard deviation of $R_T^{-0.6}$. Because the yield is proportional to $R_T^{-0.6}$ in equation (7), and the vast majority of implosions carried out at OMEGA are repeatable, the empirical correction factor in equation (7) varies by less than 10% for most OMEGA implosions. This justifies the use of this parameter to account for the occasional shot with large random nonuniformities.
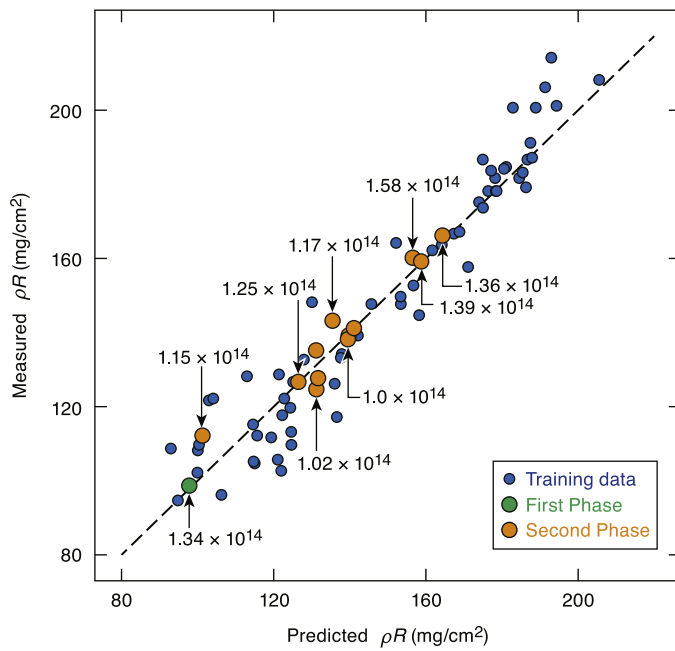
Extended Data Fig. 4 | Areal densities for high-convergence implosions are well predicted. The dashed line is the $y = x$ line on which the data points would lie if the experimental areal density were perfectly modelled by the power-law dependence. As the convergence increases, the safety margin for a nominal implosion becomes thinner and makes designing these implosions more challenging. A predictive model for the areal density was built using the framework presented in this paper (training data are shown by blue circles). The areal density was increased from that of shot 87266 (green circle) over four shot days (orange circles) until the highest values were reached in July 2018. The yields for these implosions are marked.
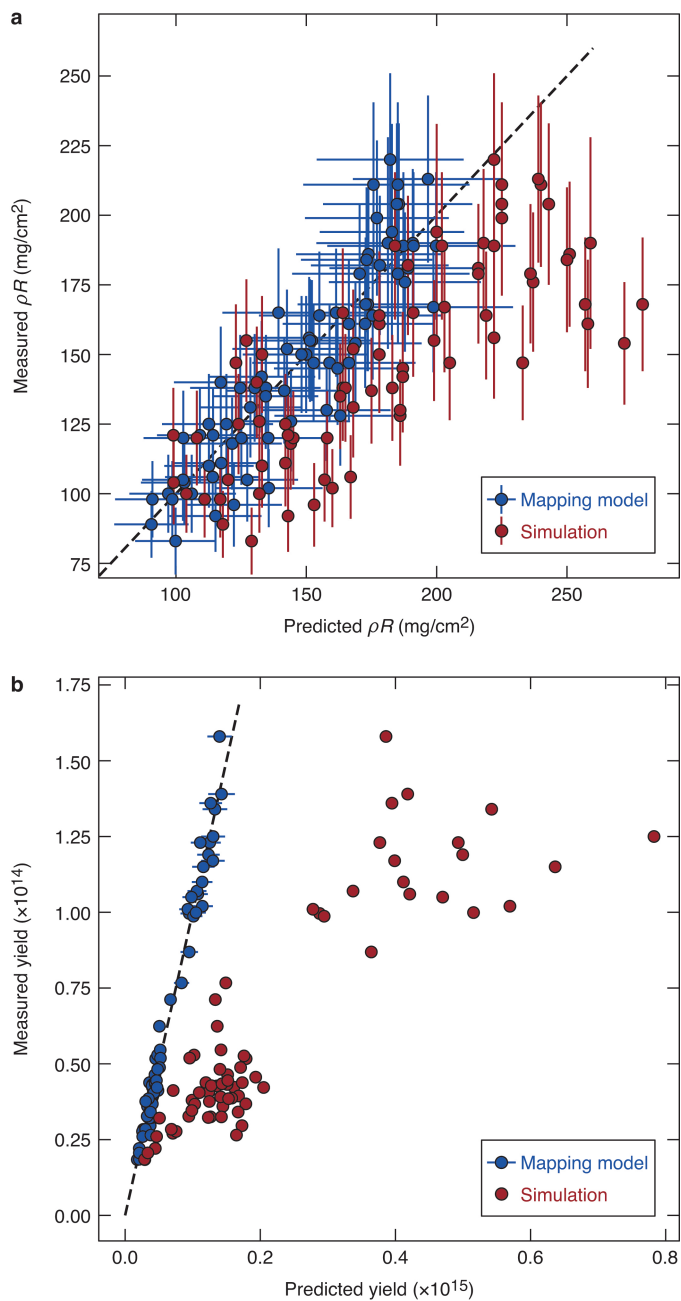
Extended Data Fig. 5 | Comparison of prediction accuracies of the statistical model and 1D simulations. a, **b**, Predictions for the areal density (**a**) and neutron yield (**b**) including the second phase of the Optimization Campaign. The predictions from the statistical model (blue) remain accurate for implosions from both phases of the Optimization Campaign, whereas the corresponding simulated 1D quantities from *LILAC* (red) remain inaccurate and overpredict potential increases in performance. Horizontal error bars and centre values represent one standard deviation and mean, respectively, for 500 draws from the posterior distribution of the statistical model. Vertical error bars represent one standard deviation for the neutron yield and areal density detectors. The dashed line is the $y = x$ line on which the data points would lie if the prediction perfectly matched the measurement.
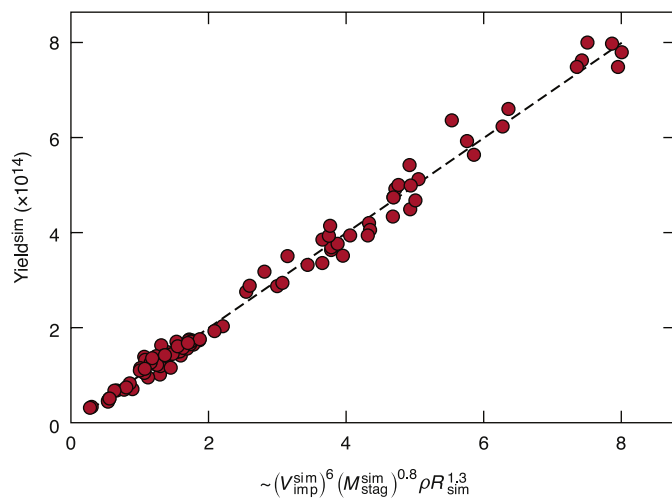
Extended Data Fig. 6 | Power-law dependence of the simulated yield on simulated parameters for initial conditions used in the Optimization Campaign. The dashed line is the $y = x$ line on which the data points would lie if the simulated yield were perfectly modelled by the power-law dependence. The simulated yield (Y^{sim}) is well represented by power-law relations to the simulated implosion velocity (V_{imp}^{sim}), mass (M_{stag}^{sim}) and areal density (ρR_{sim}). Although the exponents are not identical to those from analytical theory, physical intuition has provided a good basis for variable selection.